Two selected outcomes from the ANR DOCTOR project: NDN as microservices and an operational HTTP/NDN gateway

Xavier MARCHAL, Thibault CHOLEZ, Olivier FESTOR

LORIA (University of Lorraine, CNRS, INRIA), Nancy, France

13th June 2022

Plan

Plan

Introduction

DOCTOR Project

- DeplOyment and seCurisaTion of new functiOnalities in virtualized networking enviRonments
- Funded by ANR (French National Research Agency) : 2015-2019
- How to progressively deploy ICN in an operational state ?

Why are ICN slow to adopt ?

- Investment costs and Operational costs too expensive (CAPEX et OPEX)
- Need to have QoS guarantee

Contributions

Contributions

- Solve main **performance** and **security** problems for QoS
- Design an efficient architecture to allow **interoperability**, **deployment** and network **management**

1. Monitoring of NDN network (probe, metrics)S
2. Detection and mitigation of NDN attacks (Content Poisonning, Interest Flooding)
3. Performance evaluation of new NDN content generation
4. **Protocol and architecture to transport HTTP over NDN**
5. Ochestration of NDN Vitualized Network Functions and **Proposition of a microservice architecure**

Plan

Related work

A few application protocols adapted over ICN : VoCCN [JSB+09], HLS [XCCC12]

- But don't consider interoperability with IP

2 ways to implement interoperability

- Make every application compatible with ICN (dual-stack)
  - In the main code [QNP+15, QNT+14]
  - Via additional plugins [STC+13]
- Use *gateways*
  - Easer to develop but less efficient
  - Two implementations :
    - for HTTP, but partial [WBW+12]
    - for TCP, more flexible but limited usability[MO16]

⇒ Need for a convenient solution able to fully exploit NDN features

Conversion and naming

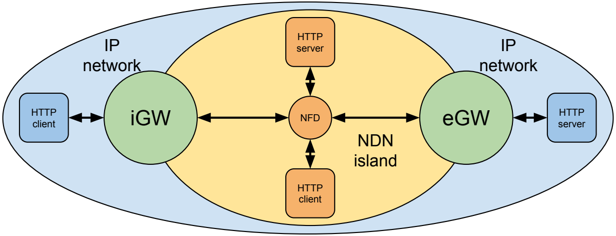

3 necessary steps to process an HTTP request (initial request, HTTP request download, answer request) :

iGW → eGW	/http/reverse_splitted_domain_name/path/client_route/sha1
eGW → iGW	/client_route/sha1(/segment)
iGW → eGW	/http/reverse_splitted_domain_name/path/sha1(/version/segment)

Hash value represents the exact way the content is requested (HTTP header options)

NDN cache hit for web contents

Reference	Reload	New user Same browser	New user Different browser
0.2%	75.8%	38.3%	0.0%

```
accept-language:
  100.0% (34415) -- en-US,en;q=0.5
accept-encoding:
  98.0% (33720) -- gzip, deflate
  2.0% (695) -- identity
accept:
  90.2% (31045) -- */*
  5.1% (1765) -- text/css,*/*;q=0.1
user-agent:
  100.0% (34415) -- Mozilla/5.0 (X11; Ubuntu; Linux x86_64; rv:50.0) Gecko/20100101 Firefox/50.0
```

```
accept-language:
  100.0% (33617) -- en-US,en;q=0.8
accept-encoding:
  100.0% (33614) -- gzip, deflate, sdch
  0.0% (3) -- gzip, deflate
accept:
  77.3% (25979) -- image/webp,image/*,*/*;q=0.8
  17.4% (5843) -- */*
user-agent:
  100.0% (33617) -- Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko)
        Chrome/55.0.2883.87 Safari/537.36
```

Maximising cache use

We consider static contents only

Smart hash computation based on selected header fields :

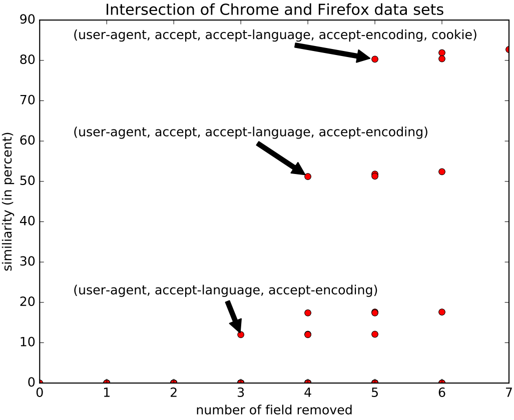

$\Rightarrow$ High level of unification by removing only a few fields

Maximising cache use

	Reference	Reload	New user Same browser	New user Different browser
No unification	0.2%	75.8%	38.3%	0.0%
Unification excluding cookie	0.2%	76.8%	37.8%	27.4%
Unification including cookie	0.2%	74.6%	78.3%	61.3%

Results :

- Different, yet similar, browser configurations can now benefit from content retrieved by others
- Improved bandwidth consumption without generating errors

Gateways performance evaluation

Median time to load a webpage

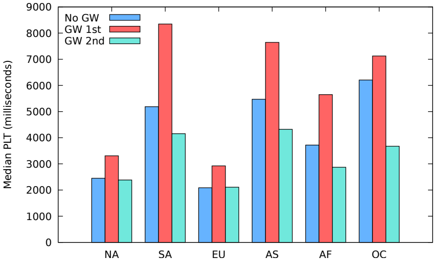

Proportion web objects well retrieved

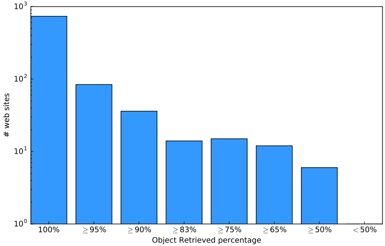

Observations :

- Similar delays, improved with caching on distant locations
- 736/1000 sites are perfectly loaded
- 966/1000 sites are loaded with more than 90% of content

Evaluation of cache hit ratio

Expérience :

- 4 Firefox puppets
- Selection of a website with a Zipf function
- Browsing time on a website follows a probabilistic model [LWD10]

Evolution of cache hit ratio in time

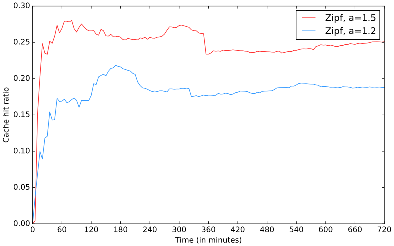

Results :

Cache hit ratio linked to popularity distribution

- 18,8% with Zipf 1,2
- 25,1% with Zipf 1,5

$\Rightarrow$ Efficient and operational transport of HTTP over NDN

Plan

Related work

NDN deployment on tweaked IP networks :

- With SDN :
  - Packets detection [VSKS13]
  - Packets labelling (CONET [SBMD+13])
  - Tunnels (NDNFlow [vAK15])
- Hybrid ICN [Net]

Deployment thanks to virtualization :

- Parallel deployment to IP (ex : [RCA+17])
- Framework : vICN (composition of functions) [SMA+17]
- Emulation : MiniCCNx [CRMa13]

⇒ Need for an architecture to ease NDN deployment and fully benefit from NFV features

Context

Network Function Virtualization (NFV) :

- Replace network equipment by VM
- Reduce CAPEX and OPEX costs
- Increased capacity to adapt and innovate

Microservices :

- Functional separation in elementary blocks
- Easier development of each service
- Better flexibility and horizontal scaling
- Each independent service consumes more resources
- Need more management
- Deployment more complex (chaining)

Microservices

Name	**Function**	**Oriented**	**Ingress/Egress cardinality**
Name Router	Route *Interest* packets	Yes	1/N
Backward Router	Route back *Data* packets	Yes	N/1
Packet Dispatcher	Split *Interest*/*Data* traffic	No	N/N
Content Store	Cache *Data* packets	No	1/1
Strategy Forwarder	Forward *Interest* packets	No	1/1 or N
Signature Verifier	Verify packets' signature	No	1/1
Name Filter	Filter on packets' name	No	1/1

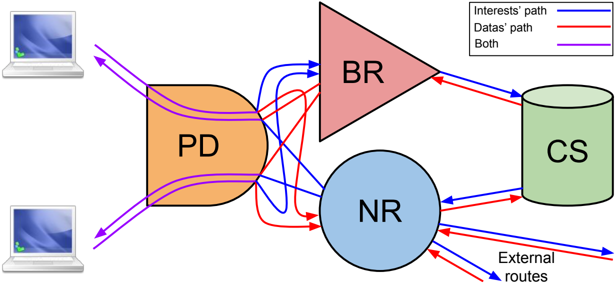

The manager

Necessary to run a microservice architecture

Mandatory operations :

- Manually or automatically deploy microservices
- Dynamic topology adaptation
- Update microservice configuration
- Scale up bottleneck components

NDN routing protocols like NLSR not mandatory inside the NDN island :

- Manager knows the whole topology
- Following events, management operations are triggered and push new configurations like an SDN controller

Evaluation of microservices' performance

Module	Throughput (Mbps)		
	Bare-Metal	Container	Difference
Name Router	2,144	1,935	-10.5%
Backward Router	1,480	1,380	-6.8%
Packet Dispatcher	2,334	2,081	-10.8%
Content Store (freshness = 0)	2,224	1,997	-10.2%
Content Store (freshness > 0)	1,151	1,010	-12,3%
Content Store (from cache)	3,431	2,852	-16,9%
Strategy Forwarder	2,281	2,058	-9.8%
Signature Verifier (RSA2048)	665	630	-5.3%
Signature Verifier (ECDSA256)	121	118	-2.5%
Name Filter	2,184	1,971	-9.8%

Observations :

- Signature verification is CPU intensive
  - Could be reduced by periodic probing
- Content Store can be a limiting element
- Use of Docker significantly decrease performance

Horizontal scaling procedure

For supporting functions

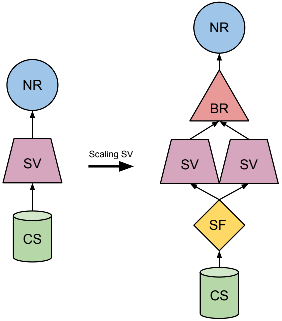

- Similar box with the same function
- BR can be replaced by another SF to remain stateless

For *Backward Router*

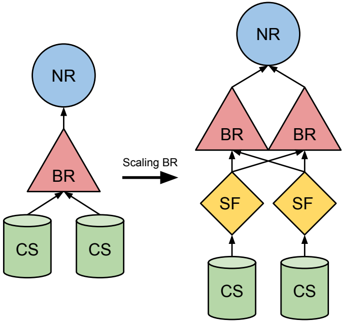

- Adding another BR would just move the bottlenexk
- Forces the next component to broadcast traffic

Scaling evaluation

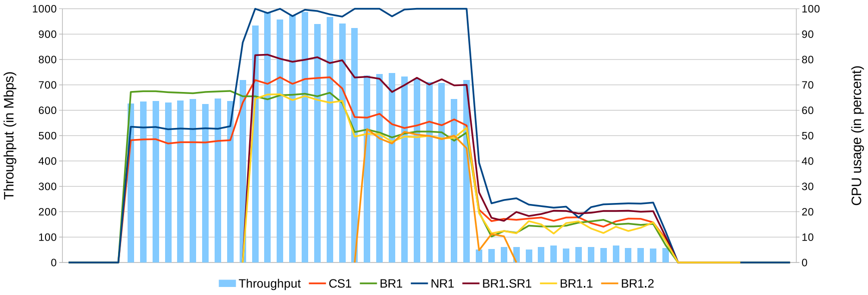

- *Backward Router* artificially limited to 67% of a CPU core
- Increased max throughput from 625 to 980 Mbps
- NR becomes the bottleneck because it needs to *broadcast* traffic to all BR instances

Orchestration of an attack mitigation

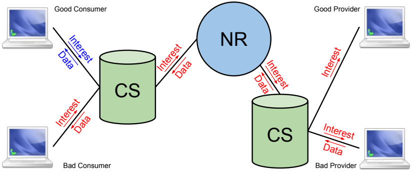

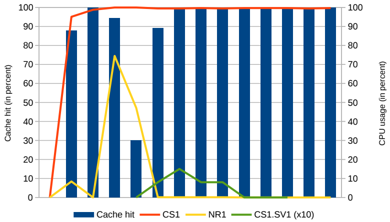

- Content Poisoning Attack
- Sudden drop of cache hit ratio triggers data packets' signature verification (deployment of SV in front of the faulty CS)
- Manager can progressively move SV toward the attack source to contain poisoned *Data*

Plan

Summary

Two selected outcomes for a better interoperability and deployment of NDN

- HTTP/NDN mapping protocol making web and NDN inter-operable, fully operational gateways
- Mutualization of web contents in NDN cache
- Architecture composed of 7 microservices and their orchestrator
- Able to dynamically react and mitigate performance and security issues

Related publications

Articles :

- **Leveraging NFV for the deployment of NDN : Application to HTTP traffic transport**, publié à IFIP/IEEE NOMS 2018
- μ**NDN : an Orchestrated Microservice Architecture for Named Data Networking**, publié à ACM ICN 2018
- **An Orchestrated NDN Virtual Infrastructure transporting Web Traffic : Design, Implementation and First Experiments with Real End-Users** publié à IEEE Communications Magazine 2019

Software contributions :

- **Both *Gateways* HTTP/NDN** (link)
- μ**NDN**, 7 orchestrated microservices for NDN (link)

… and more

Many more contributions on DOCTOR project's website :

http://www.doctor-project.org/project_outcomes.htm

Plan

References I

Carlos M.S. Cabral, Christian Esteve Rothenberg, and Maurício Ferreira Magalhães.
Mini-ccnx : Fast prototyping for named data networking.
In *Proceedings of the 3rd ACM SIGCOMM Workshop on Information-centric Networking*, ICN '13, pages 33–34, New York, NY, USA, 2013. ACM.

Van Jacobson, Diana K Smetters, Nicholas H Briggs, Michael F Plass, Paul Stewart, James D Thornton, and Rebecca L Braynard.
Voccn : voice-over content-centric networks.
In *Proceedings of the 2009 workshop on Re-architecting the internet*, pages 1–6. ACM, 2009.

Chao Liu, Ryen W White, and Susan Dumais.
Understanding web browsing behaviors through weibull analysis of dwell time.
In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 379–386. ACM, 2010.

Ilya Moiseenko and Dave Oran.
Tcp/icn : Carrying tcp over content centric and named data networks.
In *2016 conference on 3rd ACM Conference on Information-Centric Networking*, pages 112–121. ACM, 2016.

Hybrid Information-Centric Networking.
https://datatracker.ietf.org/doc/draft-muscariello-intarea-hicn/.

Xiuquan Qiao, Guoshun Nan, Yue Peng, Lei Guo, Jingwen Chen, Yunlei Sun, and Junliang Chen.
Ndnbrowser : An extended web browser for named data networking.
Journal of Network and Computer Applications, 50 :134–147, 2015.

References II

Xiuquan Qiao, Guoshun Nan, Wei Tan, Lei Guo, Junliang Chen, Wei Quan, and Yukai Tu.
Ccnxtomcat : An extended web server for content-centric networking.
Computer Networks, 75 :276–296, 2014.

R. Ravindran, A. Chakraborti, S. O. Amin, A. Azgin, and G. Wang.
5g-icn : Delivering icn services over 5g using network slicing.
IEEE Communications Magazine, 55(5) :101–107, May 2017.

Stefano Salsano, Nicola Blefari-Melazzi, Andrea Detti, Giacomo Morabito, and Luca Veltri.
Information centric networking over sdn and openflow : Architectural aspects and experiments on the ofelia testbed.
Computer Networks, 57(16) :3207–3221, 2013.

Mauro Sardara, Luca Muscariello, Jordan Augé, Marcel Enguehard, Alberto Compagno, and Giovanna Carofiglio.
Virtualized icn (vicn) : Towards a unified network virtualization framework for icn experimentation.
In *Proceedings of the 4th ACM Conference on Information-Centric Networking*, ICN '17, pages 109–115, New York, NY, USA, 2017. ACM.

Wentao Shang, Jeff Thompson, Meki Cherkaoui, Jeff Burkey, and Lixia Zhang.
Ndn. js : A javascript client library for named data networking.
In *Computer Communications Workshops (INFOCOM WKSHPS), 2013 IEEE Conference on*, pages 399–404. IEEE, 2013.

References III

Niels LM van Adrichem and Fernando A Kuipers.
Ndnflow : software-defined named data networking.
In *Network Softwarization (NetSoft), 2015 1st IEEE Conference on*, pages 1–5. IEEE, 2015.

Markus Vahlenkamp, Fabian Schneider, Dirk Kutscher, and Jan Seedorf.
Enabling information centric networking in ip networks using sdn.
In *Future Networks and Services (SDN4FNS), 2013 IEEE SDN for*, pages 1–6. IEEE, 2013.

Sen Wang, Jun Bi, Jianping Wu, Xu Yang, and Lingyuan Fan.
On adapting http protocol to content centric networking.
In *7th International Conference on Future Internet Technologies*, pages 1–6. ACM, 2012.

Hongfeng Xu, Zhen Chen, Rui Chen, and Junwei Cao.
Live streaming with content centric networking.
In *Networking and Distributed Computing (ICNDC), 2012 Third International Conference on*, pages 1–5. IEEE, 2012.

Chain of microservices equal to NFD

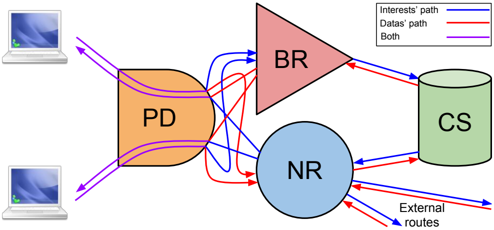

	Microservices				NFD
	PD	CS	BR	NR	
CPU core usage	100	59	89	64	100
Throughput (Mbps)	776				527
Latency (ms)	2,63				3,88

If Packet Dispatcher isnt bottleneck → 969 Mbps

Example of topology